arXiv:1610.09490v1 [stat.ML] 29 Oct 2016

A general multiblock method for structured variable selection

Tommy Löfstedt, Fouad Hadj-Selem, Vincent Guillemot,
Cathy Philippe, Nicolas Raymond, Edouard Duchesney,
Vincent Frouin and Arthur Tenenhaus

November 1, 2016

Abstract

Regularised canonical correlation analysis was recently extended to more than two sets of variables by the multiblock method Regularised generalised canonical correlation analysis (RGCCA).

Further, Sparse GCCA (SGCCA) was proposed to address the issue of variable selection. However, for technical reasons, the variable selection offered by SGCCA was restricted to a covariance link between the blocks (*i.e.*, with $\tau = 1$).

One of the main contributions of this paper is to go beyond the covariance link and to propose an extension of SGCCA for the full RGCCA model (*i.e.*, with $\tau \in [0, 1]$). In addition, we propose an extension of SGCCA that exploits structural relationships between variables within blocks. Specifically, we propose an algorithm that allows structured and sparsity-inducing penalties to be included in the RGCCA optimisation problem.

The proposed multiblock method is illustrated on a real three-block high-grade glioma data set, where the aim is to predict the location of the brain tumours, and on a simulated data set, where the aim is to illustrate the method's ability to reconstruct the true underlying weight vectors.

1 Introduction

Regularised generalised canonical correlation analysis (RGCCA) [20] is a generalisation of regularised canonical correlation analysis [23] to more than two sets of variables. RGCCA relies on a sound theoretical foundation with a well-defined optimisation criterion, while at the same time allowing the analyst to incorporate prior knowledge or hypotheses about the relationships between the blocks, as in PLS path modelling.

Sparse GCCA (SGCCA) [19] was recently proposed to address the issue of variable selection. The RGCCA criterion was modified to include ℓ_1

penalties on the outer weights vectors in order to promote sparsity. For technical reasons concerning the RGCCA algorithm, the variable selection offered by SGCCA was limited to the covariance link between blocks (*i.e.* with all $\tau_k = 1$; see below for details). One of the main contributions of this paper is to go beyond the covariance link and allow any $\tau_k \in [0, 1]$. More specifically, we present an extension of SGCCA that allows variable selection to be performed for the full RGCCA model.

The sparsity induced by the ℓ_1 penalty does not take into account any prior information on the relationships between variables within a block. Variables could for instance belong to groups, or have spatial similarities (in *e.g.* images), and could therefore benefit from the ability to perform such *structured variable selection* instead. Further, different structured penalties could be added to the different blocks, such that the regularisation adapts to the nature of the blocks. For instance, naturally related groups of variables could be selected together (or not at all), and noisy data could be constrained by spatial smoothness. We will see examples of both of these kinds of penalties below.

Therefore, we propose in this work an extension of SGCCA that allows for the exploitation of pre-given structural relationships between variables within blocks. This is achieved by introducing structured complex penalties in the model. Such penalties have recently become popular in machine learning and related fields [8] and encourage the resulting models to have a particular structure. Structured complex penalties have previously been considered in a two-block setting with canonical correlation analysis [3]. However, to combine such structured penalties with RGCCA poses new challenges for the optimisation techniques used. In this paper we propose a general multiblock algorithm that allows structured and sparsity-inducing penalties to be included in the RGCCA model.

The authors presented the main ideas behind this work in [9], and we will here give the full theoretical exposition with all details of the proposed method, including the derivation of a fast approach for projecting onto the set induced by a quadratic penalty function, and an additional example on simulated data.

2 Method

We consider several data matrices, $\mathbf{X}_1, \ldots, \mathbf{X}_K$. Each $n \times p_k$ data matrix $\mathbf{X}_k$ is called a block and represents a set of p_k variables observed on n samples. The number of variables and the nature of the variables usually differ from one block to another but the samples must be the same across the blocks. We also associate to each matrix $\mathbf{X}_k$ a column weight-vector $\mathbf{w}_k$ of dimension p_k.

Moreover, let $\mathbf{C} = (c_{kj})$ be an adjacency matrix, where $c_{kj} = 1$ if the

blocks $\mathbf{X}_k$ and $\mathbf{X}_j$ are connected, and $c_{kj} = 0$ otherwise. The main aim of RGCCA is to find block components, $\mathbf{y}_k = \mathbf{X}_k\mathbf{w}_k$, for $k = 1, \ldots, K$, that summarise the relevant information between and within the blocks, while taking into account the structural connections between blocks defined by the adjacency matrix. For that purpose, RGCCA is defined by the following optimisation problem

$$\underset{\substack{\mathbf{w}_k \in \mathbb{R}^{p_k}, \\ k=1,\ldots,K}}{\text{minimise}} \ \varphi(\mathbf{w}_1, \ldots, \mathbf{w}_K) = -\sum_{k=1}^{K}\sum_{j=1}^{K} c_{kj} g(\text{Cov}(\mathbf{X}_k\mathbf{w}_k, \mathbf{X}_j\mathbf{w}_j)), \tag{2.1}$$

$$\text{subject to} \ \ \tau_k\|\mathbf{w}_k\|_2^2 + (1-\tau_k)\text{Var}(\mathbf{X}_k\mathbf{w}_k) = 1, \tag{2.2}$$

where the constraints are defined for all $k = 1, \ldots, K$; the function g is called the inner-weighting scheme and can be any convex continuous function. Usually the function g is one of the identity, $g(x) = x$, called Horst's scheme, the absolute value, $g(x) = |x|$, called the Centroid scheme, or the square function, $g(x) = x^2$, called the Factorial scheme.

The regularisation parameters $\tau_k \in [0, 1]$ provide a way to control the trade-off between maximising correlation and maximising covariance. The above problem maximises (a function of) the covariance between connected components if $\tau_k = 1$, the correlation if $\tau_k = 0$, and a trade-off between covariance and correlation for all other values of $\tau_k \in (0, 1)$. The constraints defined by Equation 2.2 can be expressed in matrix notation as $\mathbf{w}_k^{\mathrm{T}}\mathbf{M}_k\mathbf{w}_k$ where $\mathbf{M}_k = \tau_k\mathbf{I}_{p_k} + \frac{1-\tau_k}{n-1}\mathbf{X}_k^{\mathrm{T}}\mathbf{X}_k$. We note that $\mathbf{M}_k$ is positive-semidefinite, and positive-definite when $\tau_k > 0$ or when $\mathbf{X}_k$ is of full-rank.

The SGCCA framework limits the regularisation parameters to $\tau_k = 1$ for all $k = 1, \ldots, K$. This means that variable selection is only possible for the special case with a covariance link in Equation 2.1. SGCCA is defined by the optimisation problem to

$$\begin{aligned} \underset{\substack{\mathbf{w}_k \in \mathbb{R}^{p_k}, \\ k=1,\ldots,K}}{\text{minimise}} \ & \varphi(\mathbf{w}_1, \ldots, \mathbf{w}_K) \\ \text{subject to} \ & \|\mathbf{w}_k\|_2^2 = 1, \\ & \|\mathbf{w}_k\|_1 \leq s_k, \end{aligned} \tag{2.3}$$

where both constraints are defined for all blocks $k = 1, \ldots, K$; the $\|\mathbf{w}_k\|_1 = \sum_{j=1}^{p_k} |w_{k,j}|$ is the ℓ_1-norm; the $s_k > 0$ are the radii of the ℓ_1-norm balls and determines the amount of sparsity for $\mathbf{w}_k$. The smaller s_k is, the larger the degree of sparsity for $\mathbf{w}_k$.

The ℓ_1 constraints are blind to any structure between the variables within a block and are thus not able to account for *e.g.* groups or similarities between the variables in the RGCCA model. We therefore propose to add *structured penalties* to the objective function. These structured penalties, account for such structured prior knowledge or assumptions about the variables.

The optimisation problem that we consider is thus more general, and is defined by

$$\underset{\substack{\mathbf{w}_k \in \mathbb{R}^{p_k}, \\ k=1,\dots,K}}{\text{minimise}} \quad \varphi(\mathbf{w}_1, \dots, \mathbf{w}_K) + \sum_{k=1}^{K} \omega_k \Omega_k(\mathbf{w}_k), \tag{2.4}$$

$$\begin{aligned} \text{subject to} \quad & \tau_k \|\mathbf{w}_k\|_2^2 + (1 - \tau_k)\text{Var}(\mathbf{X}_k \mathbf{w}_k) \leq 1, \\ & \|\mathbf{w}_k\|_1 \leq s_k, \end{aligned} \tag{2.5}$$

where both constraints are defined for all blocks $k = 1, \dots, K$; the functions Ω_k are the structured penalties and the Lagrange multipliers ω_k are used as regularisation parameters. The functions Ω_k are convex, but not necessarily differentiable at this point. This will be further discussed in Section 2.1.

Unfortunately, since the objective function, φ, must be convex, we must restrict the inner-weighting scheme, g, to Horst's scheme, *i.e.* to the identity $g(x) = x$.

Note that the equality in Equation 2.2 has been changed to an inequality in Equation 2.5. The reason for this is that the algorithm presented below requires the constraints to be convex. This is not really a relaxation, however, since the Karush-Kuhn-Tucker conditions require all constraints to be active at the solution, and it is always possible to find constraint parameters, s_k, such that both constraints are active for each block [25].

When the structured penalties, Ω_k, are convex, Equation 2.4 is a multiconvex function with convex constraints. This means that the function is convex with respect to one block weight vector $\mathbf{w}_k$ at the time. *I.e.* if we consider $\mathbf{w}_1, \dots, \mathbf{w}_{k-1}, \mathbf{w}_{k+1}, \dots, \mathbf{w}_K$ constant, the function is convex with respect to $\mathbf{w}_k$.

However, the structured penalties are usually neither smooth nor separable, *i.e.* they can not be written as a separable sum. We can therefore not minimise the penalties together with the smooth loss function by using smooth minimisation algorithms, and have to revert to non-smooth minimisation algorithms such as *e.g.* proximal methods. However, to compute the proximal operator of the structured penalty, we rely on separability to minimise the proximal definition coordinate-wise. Without separability, the system does not usually have an explicit solution and is therefore difficult to solve.

This means that it may be very difficult to find a minimum in the general case. Solutions exist for some particular structured penalties, but they are tailored towards a particular formulation, and can not be used for the general problem that was defined in Equation 2.4. We therefore adapt a very efficient smoothing technique proposed by Nesterov [11] to resolve both the non-smoothness and non-separability issues for a very wide and general class of structured penalties. This smoothing technique is presented in the next section.

2.1 Nesterov's smoothing technique

The structured penalties, Ω_k, considered in this paper are convex but possibly non-differentiable. The functions Ω_k must fit the framework of Nesterov, as described in [8], and will be written in the form

$$\Omega_k(\mathbf{w}_k) = \sum_{G=1}^{\mathcal{G}_k} \|\mathbf{A}_{k,G}\mathbf{w}_k\|_q = \sum_{G=1}^{\mathcal{G}_k} \max_{\|\boldsymbol{\alpha}_{k,G}\|_{q'} \leq 1} \langle \boldsymbol{\alpha}_{k,G} \,|\, \mathbf{A}_{k,G}\mathbf{w}_k \rangle,$$

in which $\mathcal{G}_k$ is the number of *groups* for the particular function Ω_k. A *group* constitutes the variables with associated non-zero entries in $\mathbf{A}_{k,G}$ (this would be *e.g.* the pixels or voxels associated with the gradient at a particular point in total variation, or a group of related variables in group $\ell_{1,q}$). The *group matrix* $\mathbf{A}_{k,G}$ is a linear operator for group G associated with the function Ω_k, for $k = 1, \ldots, K$. The function $\|\cdot\|_q$ is the standard ℓ_q-norm defined on $\mathbb{R}^p$ by $\|\mathbf{x}\|_q = \sqrt[q]{\sum_{j=1}^{p} |x_j|^q}$, with the associated dual norm $\|\cdot\|_{q'}$ for $q, q' \geq 1$. Nesterov's smoothing technique [8, 11] is formally defined as follows.

Definition 1. *Let Ω be a convex function. A sufficient condition for the application of Nesterov's smoothing technique is that Ω can be written in the form*

$$\Omega(\mathbf{w}) = \max_{\boldsymbol{\alpha} \in K} \langle \boldsymbol{\alpha} \,|\, \mathbf{A}\mathbf{w} \rangle,$$

for all $\mathbf{w} \in \mathbb{R}^p$, with K a compact convex set in a finite-dimensional vector space and $\mathbf{A}$ a linear operator between two finite-dimensional vector spaces. Given this expression for Ω, Nesterov's smoothing is defined as

$$\widehat{\Omega}(\mu, \mathbf{w}) = \langle \boldsymbol{\alpha}^* \,|\, \mathbf{A}\mathbf{w} \rangle - \frac{\mu}{2}\|\boldsymbol{\alpha}^*\|_2^2,$$

for all $\mathbf{w} \in \mathbb{R}^p$, with μ a positive real smoothing parameter and where

$$\boldsymbol{\alpha}^* = \arg\max_{\boldsymbol{\alpha} \in K} \left\{ \langle \boldsymbol{\alpha} \,|\, \mathbf{A}\mathbf{w} \rangle - \frac{\mu}{2}\|\boldsymbol{\alpha}\|_2^2 \right\}.$$

When smoothing the functions $\widehat{\Omega}_k$ this way, we obtain

$$\lim_{\mu_k \to 0} \widehat{\Omega}_k(\mu_k, \mathbf{w}_k) = \Omega_k(\mathbf{w}_k).$$

An immediate consequence is that since the functions $\widehat{\Omega}_k$ are convex and differentiable they may, for a sufficiently small value of μ_k, be used instead of Ω_k.

The gradients of the Nesterov smoothed functions, $\widehat{\Omega}_k(\mu_k, \mathbf{w}_k)$, with respect to the corresponding $\mathbf{w}_k$ are

$$\nabla_{\mathbf{w}_k} \widehat{\Omega}_k(\mu, \mathbf{w}_k) = \mathbf{A}_k^{\mathrm{T}} \boldsymbol{\alpha}_k^*.$$

The gradients are Lipschitz continuous with Lipschitz constant

$$L\big(\nabla_{\mathbf{w}_k}\widehat{\Omega}_k(\mu, \mathbf{w}_k)\big) = \frac{\|\mathbf{A}_k\|_2^2}{\mu},$$

where $\|\mathbf{A}_k\|_2$ is the spectral norm of $\mathbf{A}_k$.

2.2 Reformulation of the objective

Nesterov's smoothing technique allow us to have a smooth objective function with convex constraints. In order to find a minimiser to Equation 2.4 we must first alter the formulation slightly.

The constraints are rephrased as follows: We construct the sets $\mathcal{P}_k$ and $\mathcal{S}_k$ as $\mathcal{P}_k = \{\mathbf{x} \in \mathbb{R}^{p_k} \mid \|\mathbf{x}\|_1 \leq s_k\}$ and $\mathcal{S}_k = \{\mathbf{x} \in \mathbb{R}^{p_k} \mid \mathbf{w}_k^{\mathrm{T}}\mathbf{M}_k\mathbf{w}_k \leq 1\}$, and then form the intersection set, $\mathcal{W}_k = \{\mathbf{x} \mid \mathbf{x} \in \mathcal{P}_k \cap \mathcal{S}_k\}$. Therefore, we are interested in block weight vectors $\mathbf{w}_k$ such that $\mathbf{w}_k \in \mathcal{W}_k$. We will suppose that $\mathcal{W}_k \neq \emptyset$ and in fact assume that at least $\mathbf{0} \in \mathcal{W}_k$. We note that all $\mathcal{W}_k$ are convex sets, since $\mathcal{P}_k$ and $\mathcal{S}_k$ are convex sets.

The optimisation problem in Equation 2.4, and the final optimisation problem that we will consider in this paper, can thus be stated as

$$\begin{aligned} &\underset{\mathbf{w}_1,\ldots,\mathbf{w}_K}{\text{minimise}} \;\; \widehat{f}(\mathbf{w}_1, \ldots, \mathbf{w}_K) = \varphi(\mathbf{w}_1, \ldots, \mathbf{w}_K) + \sum_{k=1}^{K} \omega_k \widehat{\Omega}_k(\mu_k, \mathbf{w}_k) \qquad (2.6) \\ &\text{subject to} \quad \mathbf{w}_k \in \mathcal{W}_k, \quad k = 1, \ldots, K. \end{aligned}$$

This single indicative constraint for each block is equivalent to the constraints in Equation 2.4. When $\mu_k \to 0$, the two problems in Equation 2.4 and Equation 2.6 are equivalent.

The partial gradients of the objective function in Equation 2.6 with respect to each $\mathbf{w}_k$ are

$$\begin{aligned} &\nabla_{\mathbf{w}_k}\widehat{f}(\mathbf{w}_1, \ldots, \mathbf{w}_K) = \qquad (2.7) \\ &\quad - \sum_{j=1}^{K} c_{kj} g'(\mathrm{Cov}(\mathbf{X}_k\mathbf{w}_k, \mathbf{X}_j\mathbf{w}_j)) \frac{1}{n-1} \mathbf{X}_k^{\mathrm{T}}\mathbf{X}_j\mathbf{w}_j + \omega_k \mathbf{A}_k^{\mathrm{T}}\boldsymbol{\alpha}_k^*, \end{aligned}$$

for all $k = 1, \ldots, K$, where we let $\mathrm{Cov}(\mathbf{X}_k\mathbf{w}_k, \mathbf{X}_j\mathbf{w}_j) = \frac{1}{n-1}\mathbf{w}_k^{\mathrm{T}}\mathbf{X}_k^{\mathrm{T}}\mathbf{X}_j\mathbf{w}_j$, *i.e.* the unbiased sample covariance.

3 Algorithm

In order to find a solution to the problem in Equation 2.6, a multi-convex function with an indicative constraint over a convex set, we minimise it over several different parameter vectors (*i.e.* $\mathbf{w}_1, \ldots, \mathbf{w}_K$) by updating each of

the parameter vectors in turn while keeping the others fixed. *I.e.* we minimise the function over one parameter vector at the time and treat the other parameter vectors as constants during this minimisation. If each update improves the function value, gradually the function will be (locally) optimised over the entire set of parameter vectors. This principle is called block relaxation [6]. The algorithm we present in Algorithm 2 is related to the algorithm presented in [25]. However, several details need to be introduced before we discuss the proposed algorithm further.

3.1 Projection operators

We see in Algorithm 2 that orthogonal projections onto the convex sets $\mathcal{W}_k$ are required at each iteration of the algorithm, and for each block.

The projection onto the intersection of two convex sets, $\mathcal{W} = \mathcal{P} \cap \mathcal{S}$, is formulated as the unique point that minimises the problem

$$\begin{aligned} \text{proj}_{\mathcal{W}_k}(\mathbf{x}) &= \arg\min_{\mathbf{y} \in \mathcal{W}_k} \frac{1}{2}\|\mathbf{y} - \mathbf{x}\|_2^2 \\ &= \arg\min_{\mathbf{y} \in \mathbb{R}^{p_k}} \frac{1}{2}\|\mathbf{y} - \mathbf{x}\|_2^2 + \iota_{\mathcal{W}_k}(\mathbf{y}). \end{aligned} \tag{3.1}$$

where $\iota_{\mathcal{W}_k}$ is the indicator function over $\mathcal{W}_k$, *i.e.*

$$\iota_{\mathcal{W}_k}(\mathbf{x}) = \begin{cases} 0 & \text{if } \mathbf{x} \in \mathcal{W}_k, \\ \infty & \text{otherwise.} \end{cases}$$

We note that the right-most side of Equation 3.1 is in fact the proximal operator of the indicator function over the set $\mathcal{W}_k$, and thus that the proximal operator of the indicator function is the projection onto the corresponding set.

The projection onto the intersection $\mathcal{P} \cap \mathcal{S}$ can be computed using Dykstra's projection algorithm [5], stated in Algorithm 1. The sequence $(\mathbf{x}^{(s)})_{s \in \mathbb{N}}$ generated by Algorithm 1 converges to the unique point that is the solution to Equation 3.1.

We thus use Algorithm 1 to find the projection onto the intersection of the two sets $\mathcal{P}$ and $\mathcal{S}$. This is necessary in order to enforce the constraint in Equation 2.6. Three key points need to be explained in order to make Algorithm 1 clear:

(i) the projection onto $\mathcal{P}$ (Line 3),

(ii) the projection onto $\mathcal{S}$ (Line 5), and

(iii) the stopping criterion (Line 7).

These points are discussed in the following subsections.

Algorithm 1 Dykstra's projection algorithm

Require: $\mathbf{x}^{(0)}$, $\mathcal{P}$, $\mathcal{S}$, $\varepsilon > 0$
Ensure: $\mathbf{x}^{(s)} \in \mathcal{P} \cap \mathcal{S}$

$\mathbf{p}^{(0)} \leftarrow 0, \ \mathbf{q}^{(0)} \leftarrow 0$
for $s = 0, 1, 2, \ldots$ **do**
  $\mathbf{y}^{(s)} = \text{proj}_{\mathcal{P}}(\mathbf{x}^{(s)} + \mathbf{p}^{(s)})$
  $\mathbf{p}^{(s+1)} = \mathbf{x}^{(s)} + \mathbf{p}^{(s)} - \mathbf{y}^{(s)}$
  $\mathbf{x}^{(s+1)} = \text{proj}_{\mathcal{S}}(\mathbf{y}^{(s)} + \mathbf{q}^{(s)})$
  $\mathbf{q}^{(s+1)} = \mathbf{y}^{(s)} + \mathbf{q}^{(s)} - \mathbf{x}^{(s+1)}$
  if $\max\left(\|\mathbf{x}^{(s+1)} - \text{proj}_{\mathcal{P}}(\mathbf{x}^{(s+1)})\|_2, \ \|\mathbf{x}^{(s+1)} - \text{proj}_{\mathcal{S}}(\mathbf{x}^{(s+1)})\|_2\right) \leq \varepsilon$ **then**
    break
  end if
end for

3.1.1 Projection onto $\mathcal{P}$

The projection onto the ℓ_1 ball is achieved by utilising a very efficient method presented in *e.g.* [22]. This method uses the proximal operator of $\|\cdot\|_1$, the *soft thresholding operator* [12], which is defined as

$$\left(\text{prox}_{\lambda_k\|\cdot\|_1}(\mathbf{x})\right)_i = \begin{cases} x_i - \lambda_k, & \text{if } x_i > \lambda_k, \\ 0, & \text{if } |x_i| \leq \lambda_k, \\ x_i - \lambda_k, & \text{if } x_i < -\lambda_k. \end{cases} \tag{3.2}$$

This leads to the problem of finding a solution, λ_k^*, to the equation

$$\sum_{i=1}^{p_k} (|x_i| - \lambda_k^*)_+ = s_k, \tag{3.3}$$

where $(x)_+ = \max(0, x)$.

Using the parameter λ_k^* with the proximal operator results in the projection onto an ℓ_1 ball of radius s_k. *I.e.*

$$\text{proj}_{\mathcal{P}_k}(\mathbf{x}) = \text{prox}_{\lambda_k^*\|\cdot\|_1}(\mathbf{x}),$$

where thus λ_k^* is the solution of Equation 3.3.

The method we use makes the observation that if the absolute values of $\mathbf{x}$ are sorted, the solution to Equation 3.3 is found between two consecutive values of the sorted absolute x_i. The exact optimal value is then found by simply interpolating linearly (because of the nature of the Lagrange dual function of Equation 3.2) between those two values. See [22] for the details.

3.1.2 Projection onto $\mathcal{S}$

The $\mathcal{S}$ *constraint* is quadratic, which means its proximal operator is

$$\begin{aligned}\operatorname{prox}_{\lambda_k}(\mathbf{x}) &= \underset{\mathbf{y}\in\mathbb{R}^{p_k}}{\arg\min}\, \frac{1}{2}\|\mathbf{y}-\mathbf{x}\|_2^2 + \lambda_k \mathbf{y}^{\mathrm{T}}\mathbf{M}_k\mathbf{y} \\ &= \left(\mathbf{I}_{p_k} + 2\lambda_k\mathbf{M}_k\right)^{-1}\mathbf{x}.\end{aligned} \tag{3.4}$$

Let λ_k^* be the smallest λ_k such that $\mathbf{y}^{\mathrm{T}}\mathbf{M}_k\mathbf{y} \leq 1$, then

$$\operatorname{proj}_{\mathcal{S}_k}(\mathbf{x}) = \operatorname{prox}_{\lambda_k^*}(\mathbf{x}).$$

It is not feasible to compute this projection by using Equation 3.4 directly. Especially not when the number of variables is very large. *I.e.* it is not feasible to numerically find this λ_k^* directly from Equation 3.4 because of the computational effort required by the inverse. We have therefore instead devised a very efficient algorithm that rephrases the problem and then utilises the Newton-Raphson method to compute λ_k^* from a simple univariate auxiliary function that only depends on the eigenvalues of $\mathbf{M}_k$. See Appendix A for the details.

3.1.3 Stopping Criterion

Since the projection on Line 6 of Algorithm 2 is approximated (using Algorithm 1), we are actually performing an inexact projected gradient descent [17]. We must therefore make sure that the approximation is close enough that we still converge to the minimum of the objective function.

At step s of Algorithm 2, after projection onto $\mathcal{W}$ with Algorithm 1, the following inequality must be respected in order to ensure convergence to the minimum of the objective function:

$$\|\mathbf{w}^{(s+1)} - \operatorname{proj}_{\mathcal{W}}\big(\mathbf{w}^{(s+1)}\big)\|_2 < \varepsilon^{(s)},$$

where the precision, $\varepsilon^{(s)}$, must decrease like $\mathcal{O}\big(1/i_k^{4+\delta}\big)$, for any $\delta > 0$, and where i_k is the iteration counter of FISTA for block k. This follows from Proposition 2 in [17] (for FISTA, and Proposition 1 for ISTA).

Since we can not compute the distance $\|\mathbf{w}^{(s+1)} - \operatorname{proj}_{\mathcal{W}}\big(\mathbf{w}^{(s+1)}\big)\|_2$ directly (this requires a solution to the main problem we are trying to solve) and since $\mathcal{W}$ is the intersection of the convex sets $\mathcal{P}$ and $\mathcal{S}$, we may approximate it by

$$\max\Big(\|\mathbf{x}^{(s+1)} - \operatorname{proj}_{\mathcal{P}}(\mathbf{x}^{(s+1)})\|_2,\ \|\mathbf{x}^{(s+1)} - \operatorname{proj}_{\mathcal{S}}(\mathbf{x}^{(s+1)})\|_2\Big),$$

because of the well-known relation that

$$\begin{aligned}&\|\mathbf{x}^{(s+1)} - \operatorname{proj}_{\mathcal{W}}(\mathbf{x}^{(s+1)})\|_2 \\ &\quad< \kappa \cdot \max\Big(\|\mathbf{x}^{(s+1)} - \operatorname{proj}_{\mathcal{S}}(\mathbf{x}^{(s+1)})\|_2,\ \|\mathbf{x}^{(s+1)} - \operatorname{proj}_{\mathcal{P}}(\mathbf{x}^{(s+1)})\|_2\Big),\end{aligned} \tag{3.5}$$

for some positive real scalar κ.

3.2 Algorithm for Structured Variable Selection in RGCCA

We are now ready to discuss the full multiblock accelerated projected gradient method, a generalised RGCCA minimisation algorithm. This algorithm is presented in Algorithm 2.

Algorithm 2 Algorithm for structured variable selection in RGCCA

Require: $\widehat{f},\ \nabla\widehat{f},\ \mathbf{w}_k = \mathbf{w}_k^{(0)} \in \mathcal{W}_k,\ \varepsilon > 0$

Ensure: $\mathbf{w}_k^{(s)} \in \mathcal{W}_k$ such that $\varepsilon \in \partial\widehat{f}(\mathbf{w}_1^{(s)}, \ldots, \mathbf{w}_K^{(s)})$

1: **repeat**
2: **for** $k = 1$ **to** K **do**
3: $\mathbf{w}_k^{(1)} = \mathbf{w}_k^{(0)} = \mathbf{w}_k$
4: **for** $s = 1, 2, \ldots$ **do**
5: $\mathbf{y} = \mathbf{w}_k^{(s)} + \frac{k-2}{k+1}\left(\mathbf{w}_k^{(s)} - \mathbf{w}_k^{(s-1)}\right)$
6: $\mathbf{w}_k^{(s+1)} = \mathrm{proj}_{\mathcal{W}_k}\left(\mathbf{y} - t_k \nabla_{\mathbf{w}_k^{(s)}} \widehat{f}\left(\mathbf{w}_1^{(s)}, \ldots, \mathbf{y}, \ldots, \mathbf{w}_K^{(s)}\right)\right)$
7: **if** $\|\mathbf{w}_k^{(s+1)} - \mathbf{y}\|_2 \leq t_k \varepsilon$ **then**
8: **break**
9: **end if**
10: **end for**
11: $\mathbf{w}_k = \mathbf{w}_k^{(s+1)}$
12: **end for**
13: **until** $\left\|\mathbf{w}_k - \mathrm{proj}_{\mathcal{W}_k}\left(\mathbf{w}_k - t_k \nabla_{\mathbf{w}_k} \widehat{f}(\mathbf{w}_1, \ldots, \mathbf{w}_K)\right)\right\|_2 < t_k \varepsilon$, **for all** $k = 1, \ldots, K$

Any appropriate minimisation algorithm can be used in the inner-most loop of Algorithm 2. We use the fast iterative shrinkage-thresholding algorithm (FISTA) [1, 2], since it has the optimal (for first-order methods) convergence rate of $\mathcal{O}(1/s^2)$, where s is the iteration count.

FISTA requires a step size, t_k, for each block and each step of the iterative algorithm, as seen on Line 6 of Algorithm 2. If all partial gradients of the objective function in Equation 2.6, *i.e.* the gradients in Equation 2.7, are Lipschitz continuous, then we can compute the step size directly. In that case, the step sizes, t_k, are computed as the reciprocal of the sum of the Lipschitz constants of the gradients, as explained in [8]. *I.e.*, such that

$$\begin{aligned} t_k &= \left(L\left(\nabla_{\mathbf{w}_k}\widehat{f}\right)\right)^{-1} \\ &= \left(L\left(\nabla_{\mathbf{w}_k}\phi\right) + L\left(\nabla_{\mathbf{w}_k}\omega_k\widehat{\Omega}_k(\mu, \mathbf{w}_k)\right)\right)^{-1} \\ &= \left(L\left(\nabla_{\mathbf{w}_k}\phi\right) + L\left(\frac{\omega_k\|\mathbf{A}_k\|_2^2}{\mu_k}\right)\right)^{-1}, \end{aligned}$$

where μ_k is the parameter for the Nesterov smoothing; the partial gradients are from the the loss function in Equation 2.6, *i.e.* the Lipschitz constants

of the partial gradients in Equation 2.7. If some gradient is not Lipschitz continuous, or if the sum of Lipschitz constants would be zero, the step size can also be found efficiently by using backtracking line search.

Note that the main stopping criterion on Line 13 is actually performing a step of the iterative soft-thresholding algorithm (ISTA). This stopping criterion can easily be explained as follows: Assume a function on the form

$$f(\mathbf{x}) = g(\mathbf{x}) + h(\mathbf{x}),$$

where g is smooth and convex, and h is convex, but non-smooth, and whose proximal operator is known. We recall and rewrite the ISTA descent step,

$$\begin{aligned}\mathbf{x}^{(s+1)} &= \text{prox}_{th}\big(\mathbf{x}^{(s)} - t\nabla g(\mathbf{x}^{(s)})\big) \\ &= \mathbf{x}^{(s)} - tG_t(\mathbf{x}^{(s)}),\end{aligned}$$

where

$$G_t(\mathbf{x}) = \frac{1}{t}\Big(\mathbf{x} - \text{prox}_{th}\big(\mathbf{x} - t\nabla g(\mathbf{x})\big)\Big).$$

Thus, it is clear that convergence has been achieved if $G_t(\mathbf{x}^{(s)})$ is *small*. In fact, it follows from the definition of subgradients and the optimality condition of proximal operators [14] that

$$G_t(\mathbf{x}) \in \nabla g(\mathbf{x}) + \partial h\big(\mathbf{x} - tG_t(\mathbf{x})\big),$$

and that $G_t(\mathbf{x}) = \mathbf{0}$ if and only if $\mathbf{x}$ minimises $f(\mathbf{x}) = g(\mathbf{x}) + h(\mathbf{x})$.

4 Examples

We will illustrate the proposed method by two examples. The first example is on a real three-block glioma data set where the aim is to predict the location of brain tumours from gene expression (GE) and comparative genomic hybridisation (CGH) data. The second example is on a simulated data set where the aim is to see if it is possible to reconstruct the true underlying weights.

4.1 Glioma data set

We illustrate the proposed method by predicting the location of brain tumours from GE and CGH data [13]. The problem is one with three blocks: GE ($\mathbf{X}_1 \in \mathbb{R}^{53\times 15702}$), CGH ($\mathbf{X}_2 \in \mathbb{R}^{53\times 41996}$) and a dummy matrix encoding the locations ($\mathbf{X}_3 \in \mathbb{R}^{53\times 3}$). The locations were: The brain stem (DIPG), central nuclei (Midline) and supratentorial (Hemisphere).

The purpose of this example is to show the versatility of the proposed method, and to show how it can be used to build an RGCCA model with

both complex penalties and sparsity-inducing constraints, and analyse data related to the data used in [19].

The relation design was, the between-block connections, were chosen to be oriented towards prediction, and is illustrated in Figure 1. Therefore, $\mathbf{X}_1$ and $\mathbf{X}_2$ are connected to $\mathbf{X}_3$ ($c_{13} = c_{23} = 1$), but there is no connection between $\mathbf{X}_1$ and $\mathbf{X}_2$ (*i.e.*, $c_{12} = 0$). This design tends to focus on models where prediction can be made even if there is no relation between the predictor blocks, and in [19] it was indeed the case on equivalent data that this design yielded the best prediction rates among similar designs.

An ℓ_1 and a group $\ell_{1,2}$ [18, 14] constraint were added to the GE block, $\mathbf{X}_1$. An ℓ_1 together with a total variation [10], constraint were added to the CGH, $\mathbf{X}_2$, block in order to smooth the often noisy CGH data.

The regularisation constants, τ_k, for $k = 1, 2, 3$, were computed using the method of Schäfer and Strimmer [16], and were $\tau_1 = 1.0$ and $\tau_3 = 1.0$, for blocks $\mathbf{X}_1$ and $\mathbf{X}_3$, respectively. For $\mathbf{X}_2$ we were unable to compute the regularisation constant, because of the large size of the data. We therefore instead used the mean of ten regularisation constants, computed from random samples of $41\,996/2 = 20\,998$ variables each, and rounded to one decimal point. The computed mean was $\tau_2 = 0.2623\ldots \approx 0.3$.

The other constants were found by grid search with 7-fold cross-validation. The cross-validation procedure maximised the statistic

$$\begin{aligned} \mathrm{R}^2_{\mathrm{pred},\mathbf{X}_3} &= \mathrm{R}^2_{\mathrm{pred},\mathbf{X}_1\to\mathbf{X}_3} \cdot \mathrm{R}^2_{\mathrm{pred},\mathbf{X}_2\to\mathbf{X}_3} \\ &= \left(1 - \frac{\|\widehat{\mathbf{X}}_{1\to3} - \mathbf{X}_3\|_F^2}{\|\mathbf{X}_3\|_F^2}\right) \cdot \left(1 - \frac{\|\widehat{\mathbf{X}}_{2\to3} - \mathbf{X}_3\|_F^2}{\|\mathbf{X}_3\|_F^2}\right) \end{aligned} \tag{4.1}$$

in order to force a high prediction rate from both $\mathbf{X}_1$ and $\mathbf{X}_2$. *I.e.*, $\mathrm{R}^2_{\mathrm{pred},\mathbf{X}_3}$ is the combined prediction rate from the models of $\mathbf{X}_1$ and $\mathbf{X}_2$, where the product forces both blocks to predict $\mathbf{X}_3$ well. By $\widehat{\mathbf{X}}_{1\to3}$ we denote the prediction of the locations, encoded in the dummy matrix $\mathbf{X}_3$, from the

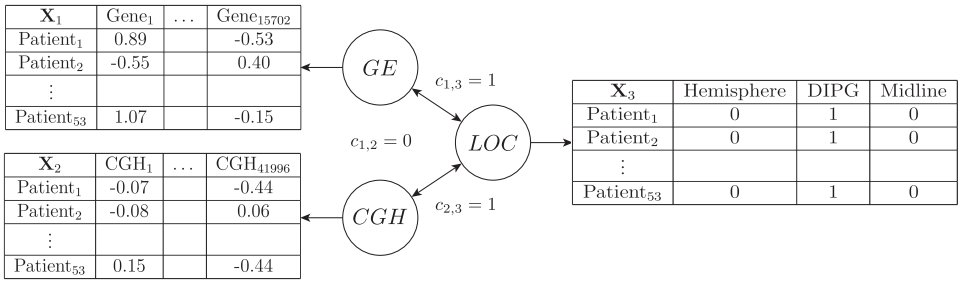

Figure 1: Path diagram for the prediction model. $\mathbf{X}_1$ (GE) is connected to $\mathbf{X}_3$ (location), $\mathbf{X}_2$ is also connected to $\mathbf{X}_3$ but $\mathbf{X}_1$ and $\mathbf{X}_2$ are not connected.

model of $\mathbf{X}_1$ and by $\widehat{\mathbf{X}}_{2\to 3}$ we denote the prediction of $\mathbf{X}_3$ from the model of $\mathbf{X}_2$. The $\|\cdot\|_F^2$ denotes the squared Frobenius norm.

The predictions were computed in the traditional way using the inner (or structural) relation [21, 15, 24], *i.e.* that the assumed relation between corresponding latent variables is linear,

$$\widehat{\mathbf{t}}_{k\to j} = \mathbf{t}_k b_{k\to j} + \boldsymbol{\varepsilon}_{k\to j},$$

where

$$b_{k\to j} = \frac{\mathbf{t}_k^{\mathrm{T}}\mathbf{t}_j}{\mathbf{t}_k^{\mathrm{T}}\mathbf{t}_k},$$

is a regression coefficient. Multiple regression is performed when there are several latent variables.

Finally, we predict with

$$\widehat{\mathbf{X}}_{k\to 3} = \mathbf{T}_k(\mathbf{T}_k^{\mathrm{T}}\mathbf{T}_k)^{-1}\mathbf{T}_k^{\mathrm{T}}\mathbf{T}_3\mathbf{W}_3^{\mathrm{T}}, = \widehat{\mathbf{T}}_{k\to 3}\mathbf{W}_3^{\mathrm{T}}$$

where $\mathbf{T}_k = [\mathbf{X}_k\mathbf{w}_{k,1}|\mathbf{X}_k\mathbf{w}_{k,2}|\cdots|\mathbf{X}_k\mathbf{w}_{k,A}]$ are the A latent variables for block k and $\mathbf{W}_3 = [\mathbf{w}_{3,1}|\cdots|\mathbf{w}_{3,A}]$ are the A weight vectors for block 3.

The criterion of the final model, as found by the grid search, was

$$\mathrm{R}^2_{\mathrm{pred},\mathbf{X}_3} = \mathrm{R}^2_{\mathrm{pred},\mathbf{X}_1\to\mathbf{X}_3} \cdot \mathrm{R}^2_{\mathrm{pred},\mathbf{X}_2\to\mathbf{X}_3} = 0.51 \cdot 0.47 \approx 0.24.$$

These numbers were computed as the means of 7-fold cross-validation.

The regularisation constant for the group $\ell_{1,2}$ penalty was thus deemed by the grid search to be $\omega_1 = 0.35$ and the ℓ_1 norm constraint had a radius of $s_1 = 13$. The regularisation constant for total variation was $\omega_2 = 0.004$, and the ℓ_1 norm constraint had a radius of $s_2 = 10.1$.

The optimisation problem that we considered in this example was thus

$$\begin{aligned}
\underset{\substack{\mathbf{w}_k\in\mathbb{R}^{p_k},\\ k=1,2,3}}{\text{minimise}}\ & \widehat{f}(\mathbf{w}_1,\mathbf{w}_2,\mathbf{w}_3) = -\mathrm{Cov}(\mathbf{X}_1\mathbf{w}_1,\mathbf{X}_3\mathbf{w}_3) - \mathrm{Cov}(\mathbf{X}_2\mathbf{w}_2,\mathbf{X}_3\mathbf{w}_3) \\
& \qquad + 0.35\cdot\widehat{\Omega}_{GL}(\mu_1,\mathbf{w}_1) + 0.004\cdot\widehat{\Omega}_{TV}(\mu_2,\mathbf{w}_2), \\
\text{subject to } & \mathbf{w}_1 \in \{\mathbf{x}\in\mathbb{R}^{p_1} \mid \|\mathbf{x}\|_1 \le 13 \wedge \|\mathbf{x}\|_2^2 \le 1\}, \\
& \mathbf{w}_2 \in \{\mathbf{x}\in\mathbb{R}^{p_2} \mid \|\mathbf{x}\|_1 \le 10.1 \wedge 0.3\cdot\|\mathbf{x}\|_2^2 + 0.7\cdot\mathrm{Var}(\mathbf{X}_2\mathbf{x}) \le 1\}, \\
& \mathbf{w}_3 \in \{\mathbf{x}\in\mathbb{R}^{p_3} \mid \|\mathbf{x}\|_2^2 \le 1\},
\end{aligned} \tag{4.2}$$

in which $\mu_1 = \mu_2 = 5\cdot 10^{-4}$. Two components were extracted using the deflation scheme [19]

$$\mathbf{X}_k \leftarrow \mathbf{X}_k - \frac{\mathbf{X}_k\mathbf{w}_k\mathbf{w}_k^{\mathrm{T}}}{\mathbf{w}_k^{\mathrm{T}}\mathbf{w}_k}.$$

4.2 Simulated data

In order to illustrate one of the main benefits of the proposed method, we performed a simulation study in which we compared the differences between the weight vectors of the proposed method and "regular" unpenalised RGCCA to the weight vectors used when generating the simulated data.

We generated data for two blocks, $\mathbf{X}_1 \in \mathbb{R}^{50\times 150}$ and $\mathbf{X}_2 \in \mathbb{R}^{50\times 100}$, defined by the models

$$\mathbf{X}_1 = \mathbf{t}_1\mathbf{w}_1^{\mathrm{T}} + \mathbf{E}_1,$$
$$\mathbf{X}_2 = \mathbf{t}_2\mathbf{w}_2^{\mathrm{T}} + \mathbf{E}_2,$$

where $\mathbf{t}_1 \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_{50})$, $\mathbf{t}_2 \sim \mathcal{N}(\mathbf{t}_1, 0.01^2\mathbf{I}_{50})$; the columns of $\mathbf{E}_1$, were random normal, $\mathbf{E}_{1,j_1} \sim \mathcal{N}(\mathbf{0}, 0.15^2\mathbf{I}_{50})$, and similarly for the columns of $\mathbf{E}_2$, such that $\mathbf{E}_{2,j_2} \sim \mathcal{N}(\mathbf{0}, 0.2^2\mathbf{I}_{50})$.

We then built an RGCCA model using Equation 2.1 and a regularised RGCCA model using Equation 2.6. The regularised RGCCA model had a total variation penalty and an ℓ_1 constraint for $\mathbf{X}_1$, and a group $\ell_{1,2}$ penalty for $\mathbf{X}_2$. The block $\mathbf{X}_2$ was not given an ℓ_1 constraint in order to illustrate the variable selection that the group $\ell_{1,2}$ penalty provides.

The groups of the group $\ell_{1,2}$ penalty were: Group 1 $= [1, \ldots, 10]$, Group 2 $= [11, \ldots, 30]$, Group 3 $= [21, \ldots, 40]$, Group 4 $= [41, \ldots, 60]$, Group 5 $= [61, \ldots, 90]$ and Group 6 $= [91, \ldots, 100]$, as illustrated in Figure 2. Note in particular that Group 2 and Group 3 are overlapping.

The regularisation constants, τ_1 and τ_2, were computed using the method of Schäfer and Strimmer [16], and were $\tau_1 = 0.33$ and $\tau_3 = 0.32$, for the blocks $\mathbf{X}_1$ and $\mathbf{X}_2$, respectively.

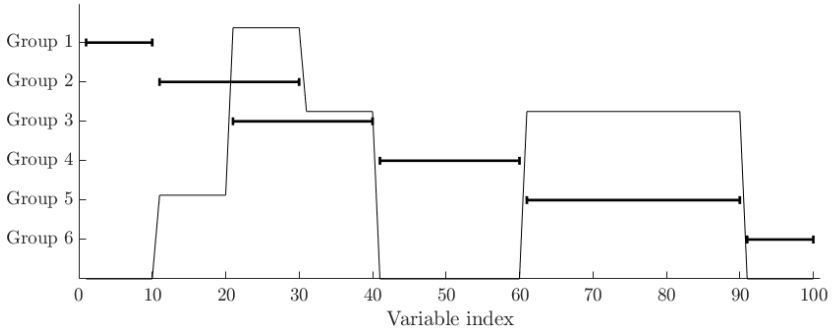

Figure 2: An illustration of the variable groups in $\mathbf{X}_2$ for the simulated data example. The horizontal thick lines correspond to a group, and the variables included in a group are indicated on the first axis. The underlying, true, weight vectors is superimposed with thin lines.

We performed a grid search with cross-validation, as in the previous example, in order to find the regularisation parameters. The constants for the regularised RGCCA model were found to be $\omega_1 = 0.61$, $s_1 = 7.7$ and $\omega_2 = 0.13$.

The optimisation problem that we considered in this example was therefore

$$\begin{aligned}
\underset{\substack{\mathbf{w}_k \in \mathbb{R}^{p_k}, \\ k=1,2}}{\text{minimise}}\ & \widehat{f}(\mathbf{w}_1, \mathbf{w}_2) = -\text{Cov}(\mathbf{X}_1\mathbf{w}_1, \mathbf{X}_2\mathbf{w}_2) \\
& \qquad + 0.61 \cdot \widehat{\Omega}_{TV}(\mu_1, \mathbf{w}_1) + 0.13 \cdot \widehat{\Omega}_{GL}(\mu_2, \mathbf{w}_2), \\
\text{subject to } & \mathbf{w}_1 \in \{\mathbf{x} \in \mathbb{R}^{p_1} \mid 0.33 \cdot \|\mathbf{x}\|_2^2 + 0.67 \cdot \text{Var}(\mathbf{X}_2\mathbf{x}) \leq 1 \wedge \|\mathbf{x}\|_1 \leq 7.7\}, \\
& \mathbf{w}_2 \in \{\mathbf{x} \in \mathbb{R}^{p_2} \mid 0.32 \cdot \|\mathbf{x}\|_2^2 + 0.68 \cdot \text{Var}(\mathbf{X}_2\mathbf{x}) \leq 1\},
\end{aligned}$$

in which $\mu_1 = \mu_2 = 5 \cdot 10^{-4}$. We only extracted one component in this example.

5 Results

5.1 Glioma data set

The locations were predicted using three approaches: From the GE data, $\mathbf{X}_1$, from the CGH data, $\mathbf{X}_2$, and from both the GE and the CGH data concatenated, *i.e.* from $\mathbf{X}_{12} = [\mathbf{X}_1 \,|\, \mathbf{X}_2]$. The GE data, $\mathbf{X}_1$, were able to predict $43/53 \approx 81$ % of the locations correctly; the CGH data, $\mathbf{X}_2$, were able to predict $38/53 \approx 72$ % of the locations correctly; and when simultaneously predicting from both the GE and the CGH data, $50/53 \approx 94$ % of the locations were correctly identified. These numbers are the means of the 7-folds of cross-validated prediction rates. These prediction rates are similar to or higher than those reported in [19] and in the present case, structure was also imposed on the weight vectors.

We performed 100 bootstrap rounds in order to assess how stable the models were, and in particular whether the models are in agreement or not between bootstrap rounds.

We computed an inter-rater agreement measure: the Fleiss' κ statistic [7], in order to asses the agreement between weight vectors computed from the 100 bootstrap rounds in terms of the selected variables for the two components. The number of times a variable is selected or not selected in the 100 bootstrap rounds is counted and summarised by the Fleiss' κ, that thus measures the agreement among the bootstrap samples. The higher the value of κ, the more stable the method is with respect to sampling; positive values means more stable than what could be expected from chance, and $\kappa = 1$ means completely stable (all samples agree entirely).

The models were stable, with all Fleiss' $\kappa > 0$. *I.e.*, the weight vectors computed in the different bootstrap rounds agreed in terms of which variables should be included or not in the model. Fleiss' κ was about 0.61 for

the first component of $\mathbf{X}_1$, and 0.35 for the second component; about 0.28 for the first component of $\mathbf{X}_2$ and 0.07 for the second component.

Figure 3 illustrates the first and second components of the models from the samples of the 100 bootstrap rounds; it is clear that the CGH data discriminates the locations very well. The coloured sections are Voronoi regions, and are meant to illustrate the location separation.

Figure 4 illustrates the first and second component of the final RGCCA model for the GE (upper left panel) and CGH (upper right panel) blocks. The lower centre panel illustrates the first components of the GE and CGH blocks plotted against each other. The separation between the locations is clear. It is also clear that the score vectors of the GE and CGH data are correlated, and that the variance of the model of the GE data is higher.

The bootstrap average number of selected variables in $\mathbf{X}_1$ was roughly 3 % in both components. In $\mathbf{X}_2$, the average number of selected variables was roughly 27 % in the first component, and roughly 40 % in the second component.

The group $\ell_{1,2}$ penalty selected 125.5 out of the 199 identified groups in the first component and 126.3 in the second component (these values are bootstrap averages). Groups were considered *strong* if they had a high ratio between the number of selected gene expressions within the group over the total number of selected gene expressions [4].

Among the top ranking groups were: Alzheimer's disease (hsa05010), which implies a relation to a supratentorial tumour (in the hemispheres) since it affects the cortex and hypocampus; "Axon guidance" (hsa04360), which implies a relation to DIPG, because of the abundance of axons in the

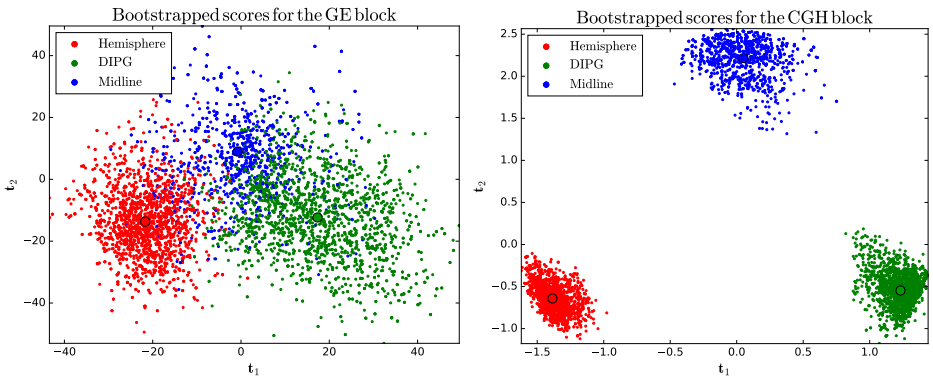

Figure 3: The first and second score vectors of the predictions of the samples from the 100 bootstrap rounds for the GE data (left) and the CGH data (right). The bootstrapped means are indicated and the Voronoi regions illustrate the separation.

brain stem; and Nucleotide excision repair (hsa03420), since it could explain tendencies towards drug resistance.

Among the groups that were excluded from the model were the Citrate cycle (TCA cycle, hsa00020). Citrate seems to be abundant in DIPG (unpublished results), but its occurrence in other locations is unknown. This would imply that it could be similarly found in the other locations or cancer types as well, and thus be a poor predictor of the tumour locations.

5.2 Simulated data

We performed cross-validation in order to find the parameters for the total variation, ℓ_1 and group lasso penalties. The obtained weight vectors were compared to the true ones (those used to generate the data) and to those of an unpenalised, or "regular", RGCCA model. The weight vectors are illustrated in Figure 5.

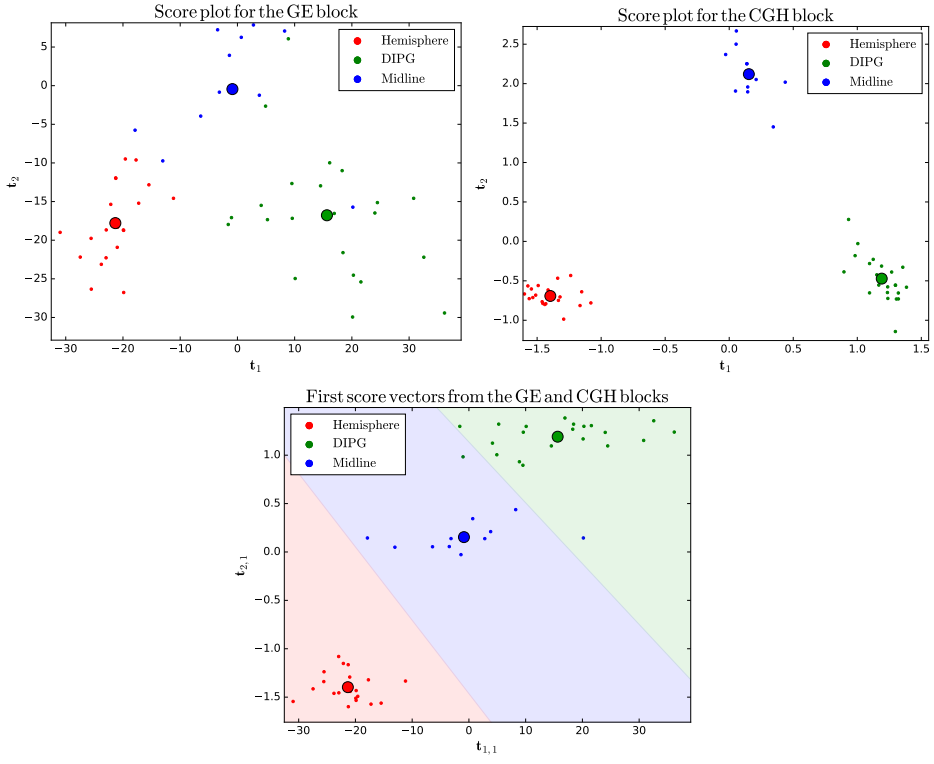

Figure 4: The first and second score vectors of the final penalised RGCCA model for the GE data (left) and the CGH data (right). The bottom plot illustrates the first score vectors plotted against each other. The predicted means are indicated in the plot by larger circles.

The weights from the unpenalised RGCCA model nicely capture the trends in the data, but are affected to a high degree by the noise.

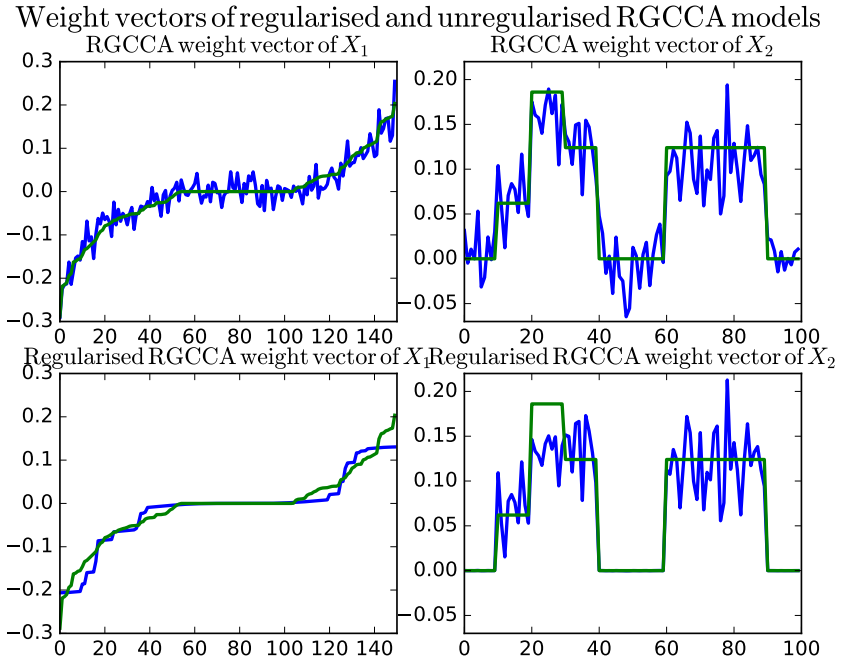
Weight vectors of regularised and unregularised RGCCA models
RGCCA weight vector of X_1
RGCCA weight vector of X_2
Regularised RGCCA weight vector of X_1
Regularised RGCCA weight vector of X_2

Figure 5: The weight vectors for an unpenalised RGCCA model (top row) for the two blocks and the weight vectors for an RGCCA model with total variation, ℓ_1 and group lasso penalties (bottom row). The total variation penalty clearly performs well in reconstructing the true underlying weight vectors (plotted in green). The group lasso penalty effectively cancels the groups with true null weights, and appears to capture the true group means fairly well, but is affected by noise. Note also that the variables with index 21 through 30 in $\mathbf{X}_2$ belongs to two groups.

As seen in Figure 5, the weight vectors obtained from the penalised model clearly performs well in reconstructing the true weights; especially so for $\mathbf{X}_1$, where the total variation penalty removes much of the noise, but leaves the true weight profile intact. The group lasso penalty manages to remove the groups with true null weights, and appears to find the mean of the other groups.

The weights of $\mathbf{X}_2$ that correspond to variable indices 20 through 30 belong to two groups (group 2 and group 3). Those should have weights that are a compromise between describing group 2 and describing group 3.

It appears they focus on describing group 2.

6 Discussion and conclusions

The proposed method solves a restriction in the SGCCA method that prevents the regularisation parameters, τ_k, for $k = 1, \ldots, K$, to take other values than one. The proposed method doesn't have this restriction, and allows the τ_k to vary between zero and one. We also showed how to extend the RGCCA and SGCCA methods by allowing complex structured penalties to be included in the model.

The proposed generalised RGCCA method was applied to gene expression and CGH data to predict the location of tumours in glioma. We used a group $\ell_{1,2}$ penalty on the GE data and a total variation penalty on the CGH data. Both data sets were also subject to an ℓ_1 constraint and the quadratic constraint used in RGCCA (a generalised norm constraint). The results are very encouraging and illustrate the importance of structured constraints.

The proposed method was also applied to simulated data were it was shown that it performed well in reconstruct the *true* weight vectors that were used in constructing the simulated data.

The authors intend to resolve the restriction that $g(x) = x$ (Horst's scheme) in Equation 2.4 in future research in order to be able to formulate even more general models. Future work also include adapting the CONESTA algorithm [8] to the present problem formulation, in order to obtain results faster, and with higher precision.

The proposed minimisation problem comprise many well-known multiblock and PLS-based methods as special cases. Examples include PLS-R, Sparse PLS, PCA, Sparse PCA, CCA, RGCCA, SGCCA, *etc.*, but the proposed method has the advantage that it allows structured and sparsity-inducing penalties to be included in the model.

7 Acknowledgements

This work was supported by grants from the French National Research Agency: ANR IA BRAINOMICS (ANR-10-BINF-04), and a European Commission grant: MESCOG (FP6 ERA-NET NEURON 01 EW1207).

References

[1] Amir Beck and Marc Teboulle. A Fast Iterative Shrinkage-Thresholding Algorithm for Linear Inverse Problems. *SIAM Journal on Imaging Sciences*, 2(1):183–202, Jan 2009.

[2] Amir Beck and Marc Teboulle. Gradient-based algorithms with applications to signal recovery problems. In Daniel P. Palomar and Yonina C. Eldar, editors, *Convex Optimization in Signal Processing and Communications*, chapter 2, pages 42–88. Cambridge University Press, 1 edition, 2009.

[3] Xi Chen and Han Liu. An Efficient Optimization Algorithm for Structured Sparse CCA, with Applications to eQTL Mapping. *Statistics in Biosciences*, 4(1):3–26, December 2011.

[4] Xi Chen, Han Liu, and Jaime G. Carbonell. Structured sparse canonical correlation analysis. volume 22 of *JMLR Workshop and Conference Proceedings*, pages 199–207, April 2012. AISTATS 2012.

[5] Patrick L. Combettes and Jean-Christophe Pesquet. Proximal splitting methods in signal processing. In H. H. Bauschke, R. S. Burachik, P. L. Combettes, V. Elser, D. R. Luke, and H. Wolkowicz, editors, *Fixed-Point Algorithms for Inverse Problems in Science and Engineering*, pages 185–212. New York: Springer, 2011.

[6] Jan De Leeuw. Block relaxation algorithms in statistics. In H. H. Bock, W. Lenski, and M. M. Richter, editors, *Information Systems and Data Analysis*, pages 308–325. Springer, Berlin, 1994.

[7] Joseph L. Fleiss. Measuring nominal scale agreement among many raters. *Psychological Bulletin*, 76(5):378–382, 1971.

[8] Fouad Hadj-Selem, Tommy Löfstedt, Vincent Frouin, Vincent Guillemot, and Edouard Duchesnay. An Iterative Smoothing Algorithm for Regression with Structured Sparsity. *arXiv:1605.09658 [stat]*, May 2016.

[9] Tommy Löfstedt, Fouad Hadj-Selem, Vincent Guillemot, Cathy Philippe, Edouard Duchesnay, Vincent Frouin, and Arthur Tenenhaus. Structured variable selection for regularized generalized canonical correlation analysis. In Hervé Abdi, Vincenzo Esposito Vinzi, Giorgio Russolillo, Gilbert Saporta, and Laura Trinchera, editors, *The Multiple Facets of Partial Least Squares Methods*, volume 173 of *Springer Proceedings in Mathematics & Statistics*. Springer, 2016.

[10] Vincent Michel, Alexandre Gramfort, Gaël Varoquaux, Evelyn Eger, and Bertrand Thirion. Total Variation Regularization for fMRI-Based Prediction of Behavior. *IEEE Transactions on Medical Imaging*, 30(7):1328–1340, 2011.

[11] Yurii Nesterov. Smooth minimization of non-smooth functions. *Mathematical Programming*, 103(1):127–152, December 2004.

[12] Neal Parikh and Stephen Boyd. *Proximal Algorithms*. Foundations and Trends in Optimization. Now Publishers Inc., 1st edition, 2013.

[13] Cathy Philippe, Stephanie Puget, Dorine A Bax, Bastien Job, Pascale Varlet, Marie-Pierre Junier, Felipe Andreiuolo, Dina Carvalho, Ricardo Reis, Lea Guerrini-Rousseau, Thomas Roujeau, Philippe Dessen, Catherine Richon, Vladimir Lazar, Gwenael Le Teuff, Christian Sainte-Rose, Birgit Geoerger, Gilles Vassal, Chris Jones, and Jacques Grill. Mesenchymal transition and PDGFRA amplification/mutation are key distinct oncogenic events in pediatric diffuse intrinsic pontine gliomas. *PloS one*, 7(2):e30313, January 2012.

[14] Zhiwei Qin, Katya Scheinberg, and Donald Goldfarb. Efficient block-coordinate descent algorithms for the Group Lasso. *Mathematical Programming Computation*, 5(2):143–169, Mar 2013.

[15] Gaston Sanchez. PLS Path Modeling with R. http://www.gastonsanchez.com/PLS_Path_Modeling_with_R.pdf, 2013.

[16] Juliane Schäfer and Korbinian Strimmer. A shrinkage approach to large-scale covariance matrix estimation and implications for functional genomics. *Statistical applications in genetics and molecular biology*, 4(1), 2005.

[17] Mark Schmidt, Nicolas Le Roux, and Francis Bach. Convergence rates of inexact proximal-gradient methods for convex optimization. arXiv:1109.2415, September 2011.

[18] Matt Silver and Giovanni Montana. Fast identification of biological pathways associated with a quantitative trait using group lasso with overlaps. *Statistical applications in genetics and molecular biology*, 11(1):Article 7, January 2012.

[19] Arthur Tenenhaus, Cathy Philippe, Vincent Guillemot, Kim-Anh Lê Cao, Jacques Grill, and Vincent Frouin. Variable Selection For Generalized Canonical Correlation Analysis. 2014.

[20] Arthur Tenenhaus and Michel Tenenhaus. Regularized Generalized Canonical Correlation Analysis. *Psychometrika*, 76(2):257–284, 2011.

[21] Michel Tenenhaus, Vincenzo Esposito Vinzi, Yves-Marie Chatelin, and Carlo Lauro. PLS path modeling. *Computational Statistics & Data Analysis*, 48:159–205, 2005.

[22] Ewout van den Berg, Mark Schmidt, Michael P. Friedlander, and Kevin Murphy. Group sparsity via linear-time projection. Technical Report TR-2008-09, Department of Computer Science, University of British Columbia, Vancouver, Canada, June 2008.

[23] Hrishikesh D. Vinod. Canonical ridge and econometrics of joint production. *Journal of Econometrics*, 4:147–166, 1976.

[24] Jacob A. Wegelin. A survey of partial least squares (pls) methods, with emphasis on the two-block case. Technical Report 371, Department of Statistics, University of Washington, Seattle, Washington, USA., March 2000.

[25] Daniela M. Witten, Robert Tibshirani, and Trevor Hastie. A penalized matrix decomposition, with applications to sparse principal components and canonical correlation analysis. *Biostatistics*, 10(3):515–534, 2009.

A The RGCCA constraint

Let $\mathbf{x}$ be a $p \times 1$ real vector and let $\mathbf{M}$ be a symmetric positive (possibly semi-) definite matrix. We consider the following optimisation problem

$$\begin{aligned} \underset{\mathbf{y}\in\mathbb{R}^p}{\text{minimise}} \;\; & \frac{1}{2}\|\mathbf{y}-\mathbf{x}\|_2^2 \\ \text{subject to} \;\; & \mathbf{y}^{\mathrm{T}}\mathbf{M}\mathbf{y} \leq c. \end{aligned} \tag{A.1}$$

This problem is equivalent to a projection of the point $\mathbf{x}$ onto a hyperellipse (a multi-dimensional ellipse) whose equation is $\mathbf{y}^{\mathrm{T}}\mathbf{M}\mathbf{y} = c$, *i.e.* a hyperellipse defined by $\mathbf{M}$ with radius c. We will thus assume that $\mathbf{x}^{\mathrm{T}}\mathbf{M}\mathbf{x} > c$, since otherwise the problem is trivial and the solution is $\mathbf{x}$. This assumption also implies that there is a single unique solution to this problem.

Let $\mathbf{X}$ be an $n \times p$ real matrix and $\sigma_1, \sigma_2, \ldots, \sigma_{\min(n,p)}$ be its singular values, then

$$\lambda_i = \begin{cases} \frac{1-\tau}{n-1}\sigma_i^2 + \tau & \text{if } i \leq \min(n,p), \\ \tau & \text{if } \min(n,p) < i \leq p, \end{cases}$$

for $i = 1, \ldots, p$, are the eigenvalues of $\mathbf{M} = \tau\mathbf{I}_p + \frac{(1-\tau)}{n-1}\mathbf{X}^{\mathrm{T}}\mathbf{X}$. Since $\mathbf{M}$ is a real symmetric matrix, there exists an orthogonal matrix $\mathbf{P}$ such that $\mathbf{P}^{-1}\mathbf{M}\mathbf{P} = \mathbf{\Lambda}$, where $\mathbf{\Lambda}$ is a diagonal matrix that contains the eigenvalues of $\mathbf{M}$.

We now define $\widetilde{\mathbf{x}} = \mathbf{P}^{-1}\mathbf{x}$ and $\widetilde{\mathbf{y}} = \mathbf{P}^{-1}\mathbf{y}$ and can thus solve a much simpler problem

$$\begin{aligned} \underset{\widetilde{\mathbf{y}}\in\mathbb{R}^p}{\text{minimise}} \;\; & \frac{1}{2}\|\widetilde{\mathbf{y}}-\widetilde{\mathbf{x}}\|_2^2 \\ \text{subject to} \;\; & \widetilde{\mathbf{y}}^{\mathrm{T}}\mathbf{\Lambda}\widetilde{\mathbf{y}} \leq c. \end{aligned}$$

The Lagrange formulation of this optimisation problem is

$$\mathcal{L}(\widetilde{\mathbf{y}}, \gamma) = \frac{1}{2}(\widetilde{\mathbf{y}}-\widetilde{\mathbf{x}})^{\mathrm{T}}(\widetilde{\mathbf{y}}-\widetilde{\mathbf{x}}) + \gamma(\widetilde{\mathbf{y}}^{\mathrm{T}}\mathbf{\Lambda}\widetilde{\mathbf{y}} - c).$$

Cancelling the gradient of the Lagrangian function $\mathcal{L}$ with respect to $\widetilde{\mathbf{y}}$ yields the following stationary equations

$$\widetilde{\mathbf{y}} - \widetilde{\mathbf{x}} + 2\gamma\mathbf{\Lambda}\widetilde{\mathbf{y}} = \mathbf{0},$$

and the solution is obtained as

$$\widetilde{\mathbf{y}} = (\mathbf{I}_p + 2\gamma\mathbf{\Lambda})^{-1}\widetilde{\mathbf{x}} = \operatorname{diag}\left(\frac{1}{1+2\gamma\lambda_1}, \ldots, \frac{1}{1+2\gamma\lambda_p}\right)\widetilde{\mathbf{x}},$$

for some γ.

Thus,

$$\widetilde{\mathbf{y}}^{\mathrm{T}}\mathbf{\Lambda}\widetilde{\mathbf{y}} = c$$

if and only if

$$\widetilde{\mathbf{x}}^{\mathrm{T}} \operatorname{diag}\left(\frac{\lambda_1}{(1+2\gamma\lambda_1)^2}, \ldots, \frac{\lambda_p}{(1+2\gamma\lambda_p)^2}\right)\widetilde{\mathbf{x}} = c.$$

Hence, we form the auxiliary function

$$f(\gamma) = \sum_{i=1}^{p} \widetilde{x}_i^2 \frac{\lambda_i}{(1+2\gamma\lambda_i)^2} - c.$$

with derivative

$$f'(\gamma) = -4\sum_{i=1}^{p} \widetilde{x}_i^2 \frac{\lambda_i^2}{(1+2\gamma\lambda_i)^3}.$$

However, we note that this approach requires us to compute all the p eigenvectors of $\mathbf{P}$, in order to find all $\widetilde{x}_i$, which may be computationally infeasible. We use the following trick to go around this problem: We note that the $p - \min(n,p)$ eigenvalues are all equal to τ, which means that

$$\sum_{i>\min(n,p)} \widetilde{x}_i^2 \frac{\lambda_i}{(1+2\gamma\lambda_i)^2} = \frac{\tau}{(1+2\gamma\tau)^2} \sum_{i>\min(n,p)} \widetilde{x}_i^2.$$

Moreover, since $\mathbf{P}$ is an orthogonal matrix,

$$\sum_{i>\min(n,p)} \widetilde{x}_i^2 = \|\mathbf{x}\|_2^2 - \sum_{i\leq\min(n,p)} \widetilde{x}_i^2,$$

and hence we implicitly know the values of $\widetilde{x}_i$ for $i > \min(n,p)$, without computing and multiplying by the corresponding eigenvectors.

We may thus rewrite the auxiliary function as

$$\begin{aligned} f(\gamma) &= \sum_{i=1}^{p} \widetilde{x}_i^2 \frac{\lambda_i}{(1+2\gamma\lambda_i)^2} - c \\ &= \frac{\tau}{(1+2\gamma\tau)^2}\left[\|\mathbf{x}\|_2^2 - \sum_{i\leq\min(n,p)} \widetilde{x}_i^2\right] + \sum_{i\leq\min(n,p)} \widetilde{x}_i^2 \frac{\lambda_i}{(1+2\gamma\lambda_i)^2} - c, \end{aligned}$$

with derivative

$$f'(\gamma) = -4\frac{\tau^2}{(1+2\gamma\tau)^3}\left[\|\mathbf{x}\|_2^2 - \sum_{i\leq\min(n,p)} \widetilde{x}_i^2\right] - 4\sum_{i\leq\min(n,p)} \widetilde{x}_i^2 \frac{\lambda_i^2}{(1+2\gamma\lambda_i)^3}.$$

Now γ can be found numerically by the Newton-Raphson method. We set an initial value of $\gamma^{(0)}$, *e.g.* $\gamma^{(0)} = 0$, and compute the sequence $(\gamma^{(s)})_s$ iteratively by

$$\gamma^{(s+1)} = \gamma^{(s)} - \frac{f(\gamma^{(s)})}{f'(\gamma^{(s)})}.$$

The algorithm stops when $|\gamma^{(s+1)} - \gamma^{(s)}| < \varepsilon$, for some small ε, *e.g.* $\varepsilon = 5 \cdot 10^{-16}$. We denote the final element of the sequence by γ^*.

We recall that the proximal operator of the quadratic function in Equation A.1 is

$$\begin{aligned}\text{prox}_\gamma(\mathbf{x}) &= \underset{\mathbf{y}\in\mathbb{R}^p}{\arg\min}\, \frac{1}{2}\|\mathbf{y}-\mathbf{x}\|_2^2 + \gamma(\mathbf{y}^{\mathrm{T}}\mathbf{M}\mathbf{y} - c) \\ &= (\mathbf{I}_p + 2\gamma\mathbf{M})^{-1}\,\mathbf{x}.\end{aligned}$$

Thus, we have the projection

$$\text{proj}_{\mathcal{S}}(\mathbf{x}) = \text{prox}_{\gamma^*}(\mathbf{x}),$$

where $\mathcal{S} = \{\mathbf{x} \in \mathbb{R}^p \mid \mathbf{x}^{\mathrm{T}}\mathbf{M}\mathbf{x} \leq c\}$. Hence, the projection is computed as

$$\mathbf{y} = (\mathbf{I}_p + 2\gamma^*\mathbf{M})^{-1}\,\mathbf{x},$$

where the inverse is computed only once, and can be computed efficiently by using the Woodbury matrix identity.